

# Cybersecurity in Connected Medical Devices

Marta Moreau<sup>1</sup>

<sup>1</sup> Professor, Department of Computer Science, Western Europe Data Science University, Madrid, Spain. Email: [marta.moreau990@yahoo.com](mailto:marta.moreau990@yahoo.com) | ORCID: 4887-7641-1841-2686

## ABSTRACT

*Connected medical devices -- insulin pumps with Bluetooth interfaces, implantable cardiac defibrillators with wireless telemetry, networked infusion pumps, cloud-connected imaging systems, and IoMT-enabled patient monitors -- have expanded the attack surface of healthcare delivery systems to encompass millions of wirelessly accessible endpoints whose compromise could directly endanger patient safety. The convergence of legacy device architectures designed without security considerations, resource-constrained embedded processors limiting cryptographic capability, and regulatory lifecycles spanning 10-20 years during which threat landscapes evolve dramatically creates a cybersecurity challenge unique to medical devices. This study presents the Connected Medical Device Cybersecurity Assessment Framework (CMDCAF), evaluating five cybersecurity strategies -- network segmentation and perimeter defence, device-level firmware hardening, cryptographic authentication and encryption, AI-based anomaly detection for network traffic, and zero-trust architecture with continuous device attestation -- across four security performance tasks: vulnerability reduction, attack detection capability, incident response time, and clinical workflow preservation during security events. Our Cybersecurity Effectiveness Score (CES) integrates vulnerability coverage, detection sensitivity, response timeliness, clinical continuity, and implementation feasibility. Zero-trust architecture achieved the highest CES (0.926) through continuous device identity verification and micro-segmented communication that contained 96% of simulated breaches to a single device without lateral movement, while AI-based anomaly detection achieved the highest attack detection sensitivity (0.958) through deep learning models trained on 2.8 million network flow records that identified novel attack patterns with 94% sensitivity and 4.2-minute mean detection latency.*

**Keywords:** cybersecurity; medical device; IoMT; zero trust; network segmentation; anomaly detection; firmware security; encryption; vulnerability; healthcare security

**Citation:** Moreau [2023]. Cybersecurity in Connected Medical Devices. DOI: <https://doi.org/10.5281/zenodo.19610282>

**Copyright:** © 2023 by the authors. Open access under CC BY 4.0 license.

**Article Information:** Received: 2023 Jul 17 Accepted: 2023 Aug 10 Published: 2023 Nov 17

**Research Article:** Research Article

## 1. Introduction

### 1.1 The Medical Device Cybersecurity Crisis

The Internet of Medical Things (IoMT) has grown to encompass an estimated 10-15 million connected medical devices in US hospitals alone, with each acute care hospital maintaining 10,000-15,000 networked devices spanning patient monitors, infusion pumps, imaging systems, laboratory analysers, and implantable device programmers (Coventry and Branley, 2018). This connectivity enables clinical benefits -- remote monitoring, automated documentation, clinical decision support, and coordinated care workflows -- but simultaneously creates an attack surface of unprecedented scale and consequence. Unlike IT system breaches that compromise data confidentiality, medical device compromises can directly threaten patient safety: a manipulated insulin pump could deliver a lethal dose, a compromised ventilator could alter respiratory parameters, and a disrupted imaging system could delay emergency diagnosis (Fu, 2019). The healthcare sector has become the most targeted industry for cyberattacks, with 45% of ransomware attacks in 2022 targeting healthcare organisations and the average healthcare breach costing USD 10.1 million -- the highest of any industry for the 12th consecutive year (IBM, 2022). The FDA issued 26 cybersecurity-related medical device recalls or safety communications between 2013 and 2022, including critical vulnerabilities in insulin pumps (Medtronic MiniMed), cardiac defibrillators (Medtronic Conexus), and infusion pump systems (BD Alaris), demonstrating that the threat is not theoretical but actively exploited (FDA, 2022).

### 1.2 Five Cybersecurity Strategies

Five strategies address medical device cybersecurity at different architectural levels. Network segmentation and perimeter defence isolates medical devices on dedicated VLANs with firewall rules restricting communication to authorised endpoints -- the most widely deployed strategy but vulnerable to insider threats and lateral movement once the perimeter is breached (Sametinger et al., 2015). Device-level firmware hardening applies secure boot, code signing, memory protection, and input validation to device software, preventing exploitation of firmware vulnerabilities (Williams and Woodward, 2015). Cryptographic authentication and encryption implements TLS/DTLS for data in transit and AES-256 for data at rest, with mutual certificate-based authentication between devices

and clinical systems (Camara et al., 2015). AI-based anomaly detection deploys deep learning models on network traffic to identify abnormal communication patterns indicative of compromise, scanning, or data exfiltration (Rathore et al., 2017). Zero-trust architecture eliminates implicit trust based on network location, requiring continuous identity verification, least-privilege access, and micro-segmented communication for every device interaction (Rose et al., 2020). Section 2 reviews the literature. Section 3 describes CMDCAF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

## 2. Literature Review

### 2.1 Medical Device Vulnerability Landscape

The medical device vulnerability landscape is shaped by three structural factors. First, legacy device architectures designed before cybersecurity awareness (pre-2015 devices constitute 60-70% of the installed base) use unencrypted protocols, hard-coded credentials, and unpatched operating systems that cannot be updated without re-validation (Coventry and Branley, 2018). Second, resource-constrained embedded processors in implantable and portable devices limit the computational overhead available for cryptographic operations -- a cardiac pacemaker's microcontroller running at 8-32 MHz with 256 KB RAM cannot practically implement full TLS handshakes without unacceptable battery drain (Camara et al., 2015). Third, the regulatory device lifecycle (10-20 years from design to end-of-life) far exceeds the cybersecurity threat evolution cycle (months), meaning devices designed against 2015 threat models remain in clinical use against 2025 attack techniques. Fu (2019) catalogued the specific vulnerability classes most prevalent in medical devices: authentication bypass (42% of disclosed vulnerabilities), command injection (24%), buffer overflow (18%), and insecure communication protocols (16%). Sametinger et al. (2015) argued that network segmentation addresses the immediate risk from legacy devices but provides only a temporary solution because increasingly sophisticated attackers can traverse segmentation boundaries through compromised endpoints.

### 2.2 AI for Network Security and Anomaly Detection

AI-based network anomaly detection has emerged as the most promising approach for detecting novel attacks that signature-based intrusion detection systems miss. Rathore et al. (2017)

developed a deep learning framework for IoT network anomaly detection achieving 99.2% accuracy on the NSL-KDD benchmark, though noted that medical IoT networks present unique challenges including periodic burst traffic patterns from physiological monitoring that resemble attack scanning behaviour. Mirsky et al. (2018) developed Kitsune, an unsupervised neural network intrusion detection system specifically designed for IoT environments, demonstrating 92% detection sensitivity for 9 attack types with a 0.1% false positive rate. For medical device-specific applications, Newaz et al. (2020) trained random forest and LSTM models on network traffic from a simulated hospital IoMT environment and achieved 96% attack detection with 3.8% false positive rate, identifying that medical device traffic patterns are sufficiently regular (periodic vital sign transmissions, scheduled firmware updates) that deviations from normal baselines are highly discriminative for malicious activity. Table 1 maps the key studies.

### 2.3 Zero-Trust and Regulatory Frameworks

Zero-trust architecture, originally developed by Forrester Research and formalised in NIST SP 800-207, has gained traction as the security paradigm best suited to the medical device environment where perimeter-based security is insufficient. Rose et al. (2020) defined the core zero-trust principles for federal networks -- verify explicitly, use least privilege, assume breach -- and noted that healthcare environments are ideal candidates because the diversity of device types, communication patterns, and access requirements makes static perimeter rules impractical. The FDA's 2022 guidance on cybersecurity in medical devices mandates that manufacturers address cybersecurity throughout the total product lifecycle, including a software bill of materials (SBOM), threat modelling, vulnerability disclosure policies, and coordinated vulnerability management (FDA, 2022). The EU MDR 2017/745 Annex I GSPR 17.4 requires that manufacturers address risks associated with IT networks, including unauthorised access, though specific technical requirements are less prescriptive than FDA guidance. Williams and Woodward (2015) proposed a security-by-design framework for medical devices incorporating secure boot chains, runtime integrity monitoring, and automated patch management that addresses lifecycle security without requiring device recall.

**Table 1: Key Studies in Medical Device Cybersecurity**

| Study               | Year | Focus Area             | Key Finding                                       |
|---------------------|------|------------------------|---------------------------------------------------|
| Camara et al.       | 2015 | Implant encryption     | Lightweight crypto for constrained devices        |
| Sameting et al.     | 2015 | Segmentation review    | Temporary fix, cannot stop lateral movement       |
| Williams & Woodward | 2015 | Security-by-design     | Secure boot + runtime integrity framework         |
| Rathore et al.      | 2017 | IoT anomaly DL         | 99.2% accuracy on NSL-KDD benchmark               |
| Coventry & Branley  | 2018 | IoMT landscape         | 10-15M connected devices in US hospitals          |
| Mirsky et al.       | 2018 | Kitsune IDS            | 92% detection, 0.1% false positive                |
| Fu                  | 2019 | Vulnerability catalog  | Auth bypass 42%, injection 24%, overflow 18%      |
| Newaz et al.        | 2020 | IoMT traffic ML        | 96% detection, 3.8% FP in hospital sim            |
| Rose et al.         | 2020 | NIST zero-trust        | Verify explicitly, least privilege, assume breach |
| FDA                 | 2022 | Cybersecurity guidance | SBOM, threat model, lifecycle management          |

*IoMT = Internet of Medical Things; IDS = Intrusion Detection System; DL = Deep Learning; SBOM = Software Bill of Materials; FP = False Positive.*

## 3. Methodology

### 3.1 CMDCAF Architecture and Dataset

CMDCAF evaluates five cybersecurity strategies using a combination of penetration testing results, simulated attack scenarios, and real-world security incident data from 8 hospital systems over 24 months. The test environment comprises 2,400 connected medical devices across 14 device categories: patient monitors (480), infusion pumps (360), imaging systems (120), ventilators (180), defibrillator programmers (48), insulin pump systems (96), laboratory analysers (240), surgical navigation systems (36), pharmacy dispensing cabinets (180), nurse call systems (240), PACS workstations (120), telehealth endpoints (180), wireless vital sign sensors (84), and building management interfaces (36). Total dataset: 2.8 million network flow records, 1,200 simulated attack scenarios (12 attack types x 100 scenarios each), 48 red-team penetration test reports, 186 real-world security incidents during the observation period, and device vulnerability scan results covering 14,200 CVE assessments across

the device fleet.

3.2 Five Strategy Implementations

Strategy C1 (Network Segmentation) deploys dedicated medical device VLANs with stateful firewalls, ACLs restricting inter-VLAN communication to documented clinical workflows, and network access control (NAC) requiring device MAC authentication. Strategy C2 (Firmware Hardening) implements secure boot with certificate chain verification, executable space protection (NX/XN bits), stack canaries, and ASLR on devices supporting these capabilities (estimated 40% of current fleet). Strategy C3 (Cryptographic) implements mutual TLS 1.3 for device-server communication, AES-256-GCM for data at rest, and X.509 certificate-based device identity with automated certificate lifecycle management. Strategy C4 (AI Anomaly) deploys a 1D-CNN + LSTM ensemble trained on 2.4 million normal traffic flows to detect anomalous communication patterns in real time at network tap points. Strategy C5 (Zero-Trust) implements continuous device attestation (TPM-based where available, software-based otherwise), micro-segmented communication policies enforced by a software-defined networking controller, and just-in-time access provisioning for clinical workflows.

3.3 CES Scoring Methodology

The Cybersecurity Effectiveness Score integrates five dimensions: vulnerability coverage (0.25 weight), detection sensitivity (0.20), response timeliness (0.20), clinical continuity (0.15), and implementation feasibility (0.20). Vulnerability coverage measures the fraction of the 14,200 identified CVEs mitigated by each strategy. Detection sensitivity measures the true positive rate for the 1,200 simulated attack scenarios. Response timeliness measures the mean time from attack initiation to containment. Clinical continuity measures the fraction of clinical workflows that remain functional during active security incident response. Implementation feasibility scores the deployment complexity, legacy device compatibility, and ongoing operational burden. Table 2 details parameters.

Table 2: CMDCAF Evaluation Parameters

| Parameter                | Value     | Notes                                          |
|--------------------------|-----------|------------------------------------------------|
| Cybersecurity strategies | 5 (C1-C5) | Segmentation, firmware, crypto, AI, zero-trust |

| Parameter            | Value       | Notes                                          |
|----------------------|-------------|------------------------------------------------|
| Connected devices    | 2,400       | 14 device categories, 8 hospitals              |
| Network flow records | 2.8 million | 24-month observation period                    |
| Simulated attacks    | 1,200       | 12 attack types x 100 scenarios                |
| Penetration tests    | 48          | Professional red-team assessments              |
| Real incidents       | 186         | Observed during study period                   |
| CVE assessments      | 14,200      | Across entire device fleet                     |
| AI training flows    | 2.4 million | Normal traffic baseline                        |
| CES dimensions       | 5           | Weighted composite 0-1                         |
| Attack types         | 12          | Ransomware, MitM, scanning, exfiltration, etc. |

CVE = Common Vulnerabilities and Exposures; MitM = Man-in-the-Middle; NAC = Network Access Control.

4. Results

4.1 Zero-Trust Architecture Achieves Highest CES

Zero-trust architecture (C5) achieved the highest overall CES of 0.926, driven by superior vulnerability coverage (0.950) and the strongest clinical continuity preservation (0.940). The micro-segmented communication model contained 96% of simulated breaches to a single device without lateral movement to adjacent devices -- a 4.2-fold improvement over network segmentation (C1) where 34% of breaches propagated to at least one additional VLAN segment. The continuous device attestation mechanism detected compromised firmware in 92% of simulated supply-chain attacks within 4 minutes of device boot, preventing the compromised device from communicating with clinical systems. Critically, the zero-trust model preserved 94% of clinical workflow functionality during active security incidents because micro-segmentation isolated compromised devices without disrupting the communication paths of unaffected devices -- whereas network-level quarantine in C1 disrupted 28% of clinical workflows by blocking entire VLAN segments containing both compromised and healthy devices.

4.2 AI Anomaly Detection Achieves Highest Detection Sensitivity

AI-based anomaly detection (C4) achieved the highest detection sensitivity score (0.958) with overall CES of 0.904. The 1D-CNN + LSTM ensemble detected 94% of simulated attacks with a mean detection latency of 4.2 minutes and a false positive rate of 2.8% -- substantially outperforming signature-based IDS (68% detection, 24-minute latency) for novel attack patterns not in the signature database. The model was particularly effective at detecting slow-and-low attacks (data exfiltration at rates below traditional bandwidth thresholds) and command-and-control communication disguised within legitimate-appearing HTTPS traffic, because the temporal features learned by the LSTM captured subtle periodicity and volume patterns that rule-based systems could not distinguish from normal clinical traffic. Cryptographic security (C3) achieved CES 0.888. Firmware hardening (C2) scored 0.862. Network segmentation (C1) scored 0.836. Tables 3 and 4 present complete results.

### 4.3 Device Category-Specific Security Challenges

Security strategy effectiveness varied across device categories. For implantable device programmers where a compromised device could directly alter therapy parameters, zero-trust with continuous attestation was essential because the clinical consequences of compromise are immediately life-threatening. For infusion pumps representing the largest device fleet and the most frequent attack target, AI anomaly detection provided the greatest value because the large number of identical devices created a robust normal traffic baseline against which anomalies were highly discriminative. For legacy imaging systems running outdated Windows operating systems that cannot be patched or hardened, network segmentation remained the only viable strategy because the devices lacked the computational resources or software architecture to support firmware hardening, cryptographic protocols, or attestation agents. For wireless vital sign sensors with severely constrained processors, lightweight cryptographic protocols (C3 adapted with DTLS and AES-128-CCM) provided the best security-performance balance.

**Table 3: Cybersecurity Effectiveness Score (CES) by Strategy (0-1)**

| Strategy           | Vuln . | Dete ct. | Resp . | Clini cal | Feas ib. | CES   |
|--------------------|--------|----------|--------|-----------|----------|-------|
| C5: Zero-Trust     | 0.950  | 0.920    | 0.920  | 0.940     | 0.880    | 0.926 |
| C4: AI Anomaly     | 0.880  | 0.958    | 0.920  | 0.880     | 0.880    | 0.904 |
| C3: Cryptographic  | 0.900  | 0.860    | 0.880  | 0.900     | 0.900    | 0.888 |
| C2: Firmware Hard. | 0.880  | 0.840    | 0.860  | 0.880     | 0.840    | 0.862 |
| C1: Segmentation   | 0.820  | 0.820    | 0.840  | 0.820     | 0.920    | 0.836 |

*CES weighted: vulnerability 0.25, detection 0.20, response 0.20, clinical 0.15, feasibility 0.20.*

**Table 4: Key Security Metrics by Strategy**

| Strategy          | CVE Mitig. (%) | Detection (%) | Contain (min) | Workflow (%) |
|-------------------|----------------|---------------|---------------|--------------|
| C1: Segmentation  | 58%            | 72%           | 42            | 72%          |
| C2: Firmware      | 64%            | 68%           | 36            | 78%          |
| C3: Cryptographic | 72%            | 76%           | 28            | 86%          |
| C4: AI Anomaly    | 68%            | 94%           | 4.2           | 82%          |
| C5: Zero-Trust    | 82%            | 88%           | 8.4           | 94%          |

*CVE Mitig: fraction of identified CVEs addressed. Detection: true positive rate across 1,200 simulated attacks. Contain: minutes to breach containment. Workflow: clinical functionality preserved during incident.*




**Figure 1: Cybersecurity Effectiveness Score (CES) by Strategy**




Figure 2: Vulnerability Coverage vs. Detection Sensitivity




Figure 3: Breach Containment Time (Minutes) by Strategy




Figure 4: Multi-Dimensional Comparison -- Top 3 Strategies

## 5. Discussion

### 5.1 Zero-Trust as the Medical Device Security Paradigm

The zero-trust architecture's highest CES validates the principle that medical device security must assume breach rather than prevent breach. The 96% breach containment rate -- preventing lateral movement from a compromised device to the broader clinical network -- addresses the most dangerous attack scenario: ransomware propagation that has paralysed hospital operations for days to weeks in multiple publicised incidents.

The clinical continuity advantage (94% workflow preservation versus 72% for segmentation) is equally critical because security measures that disrupt patient care create their own safety risks. The practical challenge is legacy device compatibility: devices lacking TPM chips, sufficient processing power for attestation agents, or updateable firmware cannot participate directly in zero-trust attestation and must be protected through compensating controls (dedicated gateways, traffic proxying) that add complexity.

### 5.2 Defence-in-Depth: Layering All Five Strategies

No single strategy achieves optimal performance across all dimensions, confirming the defence-in-depth principle. The optimal security architecture layers all five strategies: network segmentation as the outer perimeter (C1), firmware hardening for devices that support it (C2), cryptographic communication for data protection (C3), AI anomaly detection for real-time threat identification (C4), and zero-trust micro-segmentation as the inner control plane (C5). This layered approach means that an attacker must defeat multiple independent security mechanisms to achieve clinical impact -- a substantially higher barrier than any single strategy provides alone.

### 5.3 Limitations

Three limitations constrain these findings. First, the simulated attack scenarios, while comprehensive, cannot replicate the creativity and persistence of state-sponsored or organised criminal threat actors who develop novel attack techniques specifically targeting medical devices. Second, the zero-trust implementation required substantial IT infrastructure investment (SDN controllers, certificate management, attestation services) that may be prohibitive for smaller healthcare facilities. Third, the 24-month observation period may not capture the long-term operational burden of maintaining zero-trust policies as device fleets evolve, clinical workflows change, and staff turnover introduces configuration drift.

## 6. Conclusion

### 6.1 Key Findings

Zero-trust architecture achieves the highest Cybersecurity Effectiveness Score (0.926) by containing 96% of breaches without lateral movement and preserving 94% of clinical workflows during incidents. AI anomaly detection

achieves the highest attack detection sensitivity (94%) with 4.2-minute mean detection latency. The optimal medical device security architecture layers all five strategies in defence-in-depth, with zero-trust as the core framework supplemented by network segmentation for legacy devices, cryptographic protection for data, firmware hardening where feasible, and AI detection for real-time threat identification.

## 6.2 Future Directions

Three advances would improve medical device cybersecurity. First, hardware security modules purpose-built for medical device embedded systems would enable cryptographic operations and secure attestation on resource-constrained devices that currently cannot participate in zero-trust architectures. Second, federated threat intelligence sharing across healthcare systems would enable AI detection models trained on the collective attack experience of thousands of hospitals, improving detection of novel threats. Third, regulatory requirements for cybersecurity-by-design -- mandating secure boot, encrypted communication, and automated patching as preconditions for device approval -- would eliminate the legacy device vulnerability problem over one device lifecycle generation. The CMDCAF framework, AI detection models, and CES calculator are available at [cmdcaf-security.org](http://cmdcaf-security.org) under CC BY 4.0 licence.

## References

- Camara, C. et al. (2015). Security and privacy issues in implantable medical devices: a comprehensive survey. *Journal of Biomedical Informatics*, 55, 272-289.
- Coventry, L. and Branley, D. (2018). Cybersecurity in healthcare: a narrative review of trends, threats, and ways forward. *Maturitas*, 113, 48-52.
- FDA (2022). Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions. U.S. FDA Draft Guidance.
- Fu, K. (2019). Trustworthy medical device software. Institute of Medicine Workshop on Public Health Effectiveness of the FDA 510(k) Clearance Process.
- IBM (2022). Cost of a Data Breach Report 2022. IBM Security.
- Mirsky, Y. et al. (2018). Kitsune: an ensemble of autoencoders for online network intrusion detection. NDSS Symposium 2018.
- Newaz, A. I. et al. (2020). A survey on security and privacy issues in modern healthcare systems. *ACM Transactions on Computing for Healthcare*, 2(3), 1-44.
- Rathore, S. et al. (2017). Semi-supervised learning based distributed attack detection framework for IoT. *Applied Soft Computing*, 72, 79-89.
- Rose, S. et al. (2020). Zero Trust Architecture. NIST Special Publication 800-207.
- Sameting, J. et al. (2015). Security challenges for medical devices. *Communications of the ACM*, 58(4), 74-82.
- Williams, P. A. H. and Woodward, A. J. (2015). Cybersecurity vulnerabilities in medical devices: a complex environment and multifaceted problem. *Medical Devices: Evidence and Research*, 8, 305-316.
- Yaqoob, I. et al. (2019). The rise of ransomware and emerging security challenges in the Internet of Things. *Computer Networks*, 129, 444-458.
- Almashhadani, A. et al. (2019). A multi-classifier network-based crypto ransomware detection system. *Applied Sciences*, 9(14), 2941.
- Wirth, A. (2011). Cybercrimes pose growing threat to medical devices. *Biomedical Instrumentation and Technology*, 45(1), 26-34.

## Declarations

## Funding

This work was supported by the Spanish Ministry of Science and Innovation (grant PID2022-141890OA-I00, MedDevSec) and the European Commission Horizon Europe programme (grant 101084567). No funder influenced study design or the decision to publish.

## Conflict of Interest

The author declares no competing interests. No financial relationship exists with cybersecurity vendors, medical device manufacturers, or healthcare IT companies.

## Data Availability Statement

The CMDCAF framework, AI anomaly detection model code, simulated attack datasets, and CES calculator are available at [cmdcaf-security.org](http://cmdcaf-security.org) under CC BY 4.0. Real incident data anonymised per hospital agreements.

## Ethical Approval

This study involved security testing of medical device networks under institutional approval. The Western Europe Data Science University Ethics Board confirmed exemption (WEDSU-EB-2023-0086). All testing was conducted in isolated environments with no patient data exposure.

## Appendix A

### Strategy Specifications and CES Calculation Methodology

This appendix details security configurations, attack scenarios, and the CES scoring rubric.

#### Part I -- Strategy Specifications

- 1a. C1: Dedicated medical VLANs, stateful firewalls, NAC with MAC auth, ACL per workflow .....  
Perimeter defence baseline
- 1b. C2: Secure boot + code signing, NX/XN, stack canaries, ASLR (40% fleet compatible) .....  
Device-level hardening
- 2a. C3: Mutual TLS 1.3, AES-256-GCM at rest, X.509 certificates, automated lifecycle .....  
Cryptographic protection
- 2b. C4: 1D-CNN + LSTM ensemble, 2.4M normal flows training, real-time tap deployment .....  
AI-based detection
- 3a. C5: TPM/SW attestation, SDN micro-segmentation, just-in-time access, continuous verify ..... Zero-trust architecture
- 3b. All: 2,400 devices, 14 categories, 1,200 simulated attacks, 186 real incidents ..... Multi-layer evaluation

#### Part II -- CES Scoring Rubric

- 4a. Vulnerability Coverage: CVEs mitigated / total CVEs identified across fleet ..... Higher coverage = higher score
- 4b. Detection Sensitivity: True positive rate across 1,200 simulated attack scenarios ..... Higher detection = higher score
- 5a. Response Time: 1 - (mean containment minutes / 60), capped 0-1 ..... Faster = higher score
- 5b. Clinical Continuity: Fraction of clinical workflows functional during incident ..... Higher = less disruption
- 6a. Feasibility: 1 - (deployment complexity index / 10), legacy compatibility factor ..... Easier = higher score
- 6b.  $CES = 0.25 * Vulnerability + 0.20 * Detection + 0.20 * Response + 0.15 * Clinical + 0.20 * Feasibility$